

DISCRIMINANT ANALYSIS WITH ADAPTIVELY POOLED COVARIANCE

NOAH SIMON AND ROB TIBSHIRANI

ABSTRACT. Linear and Quadratic Discriminant analysis (LDA/QDA) are common tools for classification problems. For these methods we assume observations are normally distributed within group. We estimate a mean and covariance matrix for each group and classify using Bayes theorem. With LDA, we estimate a single, pooled covariance matrix, while for QDA we estimate a separate covariance matrix for each group. Rarely do we believe in a homogeneous covariance structure between groups, but often there is insufficient data to separately estimate covariance matrices. We propose ℓ_1 -PDA, a regularized model which adaptively pools elements of the precision matrices. Adaptively pooling these matrices decreases the variance of our estimates (as in LDA), without overly biasing them. In this paper, we propose and discuss this method, give an efficient algorithm to fit it for moderate sized problems, and show its efficacy on real and simulated datasets.

Keywords: Lasso, Penalized, Discriminant Analysis, Interactions, Classification

1. INTRODUCTION

Consider the usual two class problem: our data consists of n observations, each observation with a known class label $\in \{1, 2\}$, and p covariates measured per observation. Let y denote the n -vector corresponding to class (with n_1 observations in class 1 and n_2 in class 2), and X , the n by p matrix of covariates. We would like to use this information to classify future observations.

We further assume that, given class $y(l)$, each observation, x_l , is independently normally distributed with some class specific mean $\mu_{y(l)} \in \mathbb{R}^p$ and covariance $\Sigma_{y(l)}$, and that $y(l)$ has prior probability π_1 of coming from class 1 and π_2 from class 2. From here we estimate the two mean vectors, covariance matrices, and prior probabilities and use these estimates with Bayes theorem to classify future observations. In the past a number of different methods have been proposed to estimate these parameters. The simplest is Quadratic Discriminant Analysis (QDA)

which estimates the parameters by their maximum likelihood estimates

$$\pi_k = \frac{n_k}{n}$$

$$\hat{\mu}_k = \frac{1}{n_k} \sum_{y(l)=k} x_l$$

and

$$\hat{\Sigma}_k = \frac{1}{n_k} \sum_{y(l)=k} (x_l - \mu_k) (x_l - \mu_k)^\top.$$

To classify a new observation x , one finds the class with the highest posterior probability. This is equivalent in the two class case to considering

$$D(x) = \log \left(\frac{\pi_1}{\pi_2} \right) - \frac{1}{2} (x - \mu_1)^\top \Sigma_1^{-1} (x - \mu_1) \\ + \frac{1}{2} (x - \mu_2)^\top \Sigma_2^{-1} (x - \mu_2) + \log \det \left(\Sigma_1^{-1/2} \Sigma_2^{1/2} \right)$$

and if $D(x) > 0$ then classifying to class 2, otherwise to class 1.

Linear Discriminant Analysis (LDA) is a similar but more commonly used method. It estimates the parameters by a restricted MLE — the covariance matrices in both classes are constrained to be equal. So, for LDA

$$\hat{\Sigma}_1 = \hat{\Sigma}_2 = \frac{1}{n} \sum_{l=1}^n (x_l - \mu_{y(l)}) (x_l - \mu_{y(l)})^\top$$

While one rarely believes that the covariance matrices are exactly equal, often the decreased variance from pooling the estimates greatly outweighs the increased bias.

Friedman [1989] proposed Regularized Discriminant Analysis (RDA) noting that one can partially pool the covariance matrices and find a more optimal bias/variance tradeoff. He estimates Σ_k by a convex combination of the LDA and QDA estimates

$$\hat{\Sigma}_k = \lambda \hat{\Sigma}_k^{\text{QDA}} + (1 - \lambda) \hat{\Sigma}_k^{\text{LDA}}$$

where λ is generally determined by cross-validation.

We extend the idea of partially pooling the covariance matrices in a different direction. We make the further assumption that for most i, j , $(\Sigma_1^{-1})_{i,j} \approx (\Sigma_2^{-1})_{i,j}$; that the partial covariance matrices are mostly

element-wise equal (or nearly equal). Intuitively this says that conditional on all other variables, most pairs of covariates interact identically in both groups.

Given this assumption, the natural approach is to find a restricted MLE where the number of non-zero entries in $\Sigma_1^{-1} - \Sigma_2^{-1}$ is constrained to be less than some c . ie. to find

$$\begin{aligned} \text{argmax} \quad & \ell_1(\mu_1, \Sigma_1) + \ell_2(\mu_2, \Sigma_2) \\ \text{s.t.} \quad & \|\Sigma_1^{-1} - \Sigma_2^{-1}\|_0 \leq c \\ & \Sigma_1, \Sigma_2 \text{ Positive Semi-Definite} \end{aligned}$$

where ℓ_k is the Gaussian log likelihood of the observations in class k ,

$$\ell_k(\mu_k, \Sigma_k) = -\frac{n_k}{2} \log(2\pi) + \frac{n_k}{2} \log \det(\Sigma_k^{-1}) + \sum_{y(l)=k} (x_l - \mu_k)^\top \Sigma_k^{-1} (x_l - \mu_k)$$

and $\|\cdot\|_0$ is the number of nonzero elements. Unfortunately, this problem is not convex and would require a combinatorial search. Instead we consider a convex relaxation

$$\begin{aligned} (1) \quad & \text{argmax} \quad \ell_1(\mu_1, \Sigma_1) + \ell_2(\mu_2, \Sigma_2) \\ (2) \quad & \text{s.t.} \quad \|\Sigma_1^{-1} - \Sigma_2^{-1}\|_1 \leq c \\ (3) \quad & \Sigma_1, \Sigma_2 \text{ Positive Semi-Definite} \end{aligned}$$

where $\|\cdot\|_1$ is the sum of the absolute value of the entries. Because the $|\cdot|$ is not differentiable at 0, solutions to (1) have few nonzero entries in $\Sigma_1^{-1} - \Sigma_2^{-1}$ with the sparsity level dependent on c . There is a large literature about using ℓ_1 penalties to promote sparsity (Tibshirani [1996], Chen et al. [1998], among others), and in particular sparsity has been applied in a similar framework for graphical models [Banerjee et al., 2008]. Also recently, a very similar model to that which we propose has been applied to joint estimation of partial dependence among many similar graphs [Danaher et al., 2011]. The astute reader may note that (1) is not jointly convex in μ and Σ^{-1} . However, we can still find the global maximum — for fixed μ_1 and μ_2 it is convex, and, as we later show, our estimates of μ_1 and μ_2 are completely independent of our estimates of Σ_1 , and Σ_2 .

The problem (1) has an equivalent Lagrangian form (which we will write as a minimization for future convenience)

$$\begin{aligned} (4) \quad & \text{argmin} \quad -\ell_1(\mu_1, \Sigma_1) - \ell_2(\mu_2, \Sigma_2) + \lambda \|\Sigma_1^{-1} - \Sigma_2^{-1}\|_1 \\ (5) \quad & \text{s.t.} \quad \Sigma_1, \Sigma_2 \text{ Positive Semi-Definite} \end{aligned}$$

This is the objective which we will focus on in this paper. We will call its solution “ ℓ_1 Pooled Discriminant Analysis” (ℓ_1 -PDA). For $\lambda = 0$ these are just QDA estimates and for λ sufficiently large, just LDA estimates.

In this paper, we examine the ℓ_1 -PDA objective; we discuss the connections between ℓ_1 -PDA and estimating interactions in a logistic model; we show the efficacy of ℓ_1 -PDA on real and simulated data; and we give an efficient algorithm to fit ℓ_1 -PDA based on the alternating direction method of moments (ADMM).

1.1. Reductions. One may note that our objective (4) is not jointly convex in μ_k and Σ_k , however this is not a problem (the optimization splits nicely). For a fixed Σ_1 , μ_1^* minimizes

$$\frac{1}{2} \sum_{y(l)=1} (x_l - \mu_1)^\top \Sigma_1^{-1} (x_l - \mu_1).$$

This is true iff

$$\Sigma_1^{-1} \sum_{y(l)=1} (x_l - \mu_1^*) = \underline{0}.$$

Thus, $\mu_1^* = \bar{x}_1 = \frac{1}{n_1} \sum_{y(l)=1} x_l$ is the sample mean from class 1, and similarly μ_2^* is the sample mean from class 2. We can simplify our objective (4) by substituting the sample means in for μ_1^* and μ_2^* and noting that

$$\begin{aligned} \sum_{y(l)=1} (x_l - \bar{x}_1)^\top \Sigma_1^{-1} (x_l - \bar{x}_1) &= \frac{1}{2} \sum_{y(l)=1} \text{tr} \left[(x_l - \bar{x}_1)^\top \Sigma_1^{-1} (x_l - \bar{x}_1) \right] \\ &= n_1 \sum_{y(l)=1} \text{tr} \left[\Sigma_1^{-1} (x_l - \bar{x}_1) (x_l - \bar{x}_1)^\top / n_1 \right] \\ &= n_1 \text{tr} \left[\Sigma_1^{-1} \sum_{y(l)=1} (x_l - \bar{x}_1) (x_l - \bar{x}_1)^\top / n_1 \right] \\ &= n_1 \text{tr} [\Sigma_1^{-1} S_1]. \end{aligned}$$

where $\hat{\Sigma}_1$ is the sample covariance matrix for class 1.

Our new objective is

$$(6) \quad \min_{\Sigma_1, \Sigma_2} -n_1 \log \det(\Sigma_1^{-1}) + n_1 \text{tr}(\Sigma_1^{-1} S_1) - n_2 \log \det(\Sigma_2^{-1})$$

$$(7) \quad + n_2 \text{tr}(\Sigma_2^{-1} S_2) + \lambda \|\Sigma_1^{-1} - \Sigma_2^{-1}\|_1$$

subject to Σ_1 and Σ_2 positive semi-definite (PSD). This is a jointly convex problem in Σ_1^{-1} and Σ_2^{-1} .

2. SOLUTION PROPERTIES

There is a vast literature on using ℓ_1 norms to induce sparsity. In this section we will inspect the optimality conditions for our particular problem to gain some insight. We begin by reparametrizing objective (17) in terms of $\Delta = (\Sigma_1^{-1} - \Sigma_2^{-1})/2$, and $\Theta = (\Sigma_1^{-1} + \Sigma_2^{-1})/2$

$$(8) \quad \min_{\Delta, \Theta} -n_1 \log \det(\Delta + \Theta) + n_1 \operatorname{tr}([\Delta + \Theta] S_1) - n_2 \log \det(\Theta - \Delta) \\ (9) \quad + n_2 \operatorname{tr}([\Theta - \Delta] S_2) + \lambda \|\Delta\|_1$$

To find the Karush-Kuhn optimality conditions, we take the subgradient of this expression and set it equal to 0. We see that

$$(10) \quad -n_1 \left(\hat{\Delta} + \hat{\Theta} \right)^{-1} + n_1 S_1 - n_2 \left(\hat{\Theta} - \hat{\Delta} \right)^{-1} + n_2 S_2 + \lambda \partial(\hat{\Delta}) = 0$$

and

$$(11) \quad -n_1 \left(\hat{\Delta} + \hat{\Theta} \right)^{-1} + n_1 S_1 + n_2 \left(\hat{\Delta} - \hat{\Theta} \right)^{-1} - n_2 S_2 = 0$$

where $\hat{\Delta}$ and $\hat{\Theta}$ minimize the objective and $\partial(\Delta)$ is a p by p matrix with

$$\partial(\Delta)_{i,j} = \begin{cases} \operatorname{sign}(\Delta)_{i,j}, & \text{if } \Delta_{i,j} \neq 0 \\ \in [-1, 1], & \text{if } \Delta_{i,j} = 0 \end{cases}$$

Now, we can substitute Σ_1^{-1} and Σ_2^{-1} back in to the subgradient equations:

$$(12) \quad n_1 \left(S_1 - \hat{\Sigma}_1 \right) - n_2 \left(S_2 - \hat{\Sigma}_2 \right) + \lambda \partial(\hat{\Sigma}_1^{-1} - \hat{\Sigma}_2^{-1}) = 0$$

and

$$(13) \quad S_{\text{pool}} \equiv \frac{n_1 S_1 + n_2 S_2}{n_1 + n_2} = \frac{n_1 \hat{\Sigma}_1 + n_2 \hat{\Sigma}_2}{n_1 + n_2}.$$

We find these optimality conditions curious as they directly involve $\hat{\Sigma}_k$ rather than $\hat{\Sigma}_k^{-1}$. Equation (12) shows that the solution will have a sparse difference $\hat{\Sigma}_1^{-1} - \hat{\Sigma}_2^{-1}$. Though somewhat unintuitive, it parallels the KKT conditions for the Lasso and other ℓ_1 penalized problems. In particular, because the subgradient of $\|\Delta\|_1$ can take a variety of values for $\Delta_{i,j} = 0$, the optimality conditions are often satisfied with $\Delta_{i,j} = 0$ for many i, j . Equation (13) shows us that the pooled average of our estimates is unchanged ($S_{\text{pool}} = \hat{\Sigma}_{\text{pool}}$). Given the form of our penalty we find it interesting that the pooled average of the $\hat{\Sigma}_k$ is constant (independent of λ) rather than some convex combination of the $\hat{\Sigma}_k^{-1}$.

From these optimality conditions one can easily find the optimal solutions at both ends of our path (for $\lambda = 0$ and λ sufficiently large). If S_1 and S_2 are full rank, then for $\lambda = 0$ the optimality conditions are satisfied by the QDA solution with $\partial = 0$, and for $\lambda > \lambda_{\text{extrmax}} \equiv n_1 n_2 \|S_1 - S_2\|_\infty / (n_1 + n_2)$ the conditions are satisfied by the LDA solution with $\partial = n_1 n_2 (S_1 - S_2) / [\lambda (n_1 + n_2)]$. In Section 5, we give a pathwise algorithm to fit ℓ_1 -PDA along our path of λ -values from λ_{max} to 0.

2.1. When is the problem ill posed? Recall that if S_1 or S_2 is not full rank, then the QDA solution is undefined. In our case one can see that as $\lambda \rightarrow 0$ we still have this difficulty, however for $\lambda > 0$, so long as $S_{\text{pool}} = (n_1 S_1 + n_2 S_2) / (n_1 + n_2)$ is full rank, our solution is well defined. In the case that S_{pool} is not full rank, then the solution is ill-defined for all λ .

3. FORWARD VS BACKWARD MODEL

So far we have assumed a model in which the x -values are generated given the class assignments. We will henceforth refer to this as the “backward generative model” or backward model. Many other approaches to classification use a “forward generative model” wherein we consider the class assignments to be generated from the x -values (eg. logistic regression). Our backward model has a corresponding forward model. By Bayes theorem we have

$$\begin{aligned} P(y = 1|x) &= \frac{\pi_1 \exp(l_1)}{\pi_2 \exp(l_2) + \pi_1 \exp(l_1)} \\ &= \frac{\exp[\log(\pi_1/\pi_2) + l_1 - l_2]}{1 + \exp[\log(\pi_1/\pi_2) + l_1 - l_2]} \end{aligned}$$

where

$$l_k = -(x - \hat{\mu}_k)^\top \hat{\Sigma}_k^{-1} (x - \hat{\mu}_k) / 2.$$

We can simplify this to get a better handle on it. Some algebra gives us

$$(14) \quad \text{logit}[P(y = 1|x)] = \log(\pi_1/\pi_2) + \mu_2^\top \Sigma_2^{-1} \mu_2 / 2 - \mu_1^\top \Sigma_1^{-1} \mu_1 / 2$$

$$(15) \quad + (\Sigma_1^{-1} \mu_1 - \Sigma_2^{-1} \mu_2)^\top x + x^\top (\Sigma_2^{-1} - \Sigma_1^{-1}) x / 2.$$

where $\text{logit}(p) = p/(1-p)$. This is just a logistic model with interactions and quadratic terms. In general a logistic model takes the form

$$\text{logit}[P(y = 1|x)] = \beta_0 + \sum \beta_i x_i + \sum_{i \leq j} \gamma_{i,j} x_i x_j$$

or in matrix form

$$(16) \quad \text{logit} [P(y = 1|x)] = \beta_0 + \beta^\top x + x^\top \Gamma x/2$$

So our forward generative model in (14) is a logistic model with

$$\begin{aligned} \beta_0 &= \log(\pi_1/\pi_2) + \mu_2^\top \Sigma_2^{-1} \mu_2/2 - \mu_1^\top \Sigma_1^{-1} \mu_1/2 \\ \beta &= \Sigma_1^{-1} \mu_1 - \Sigma_2^{-1} \mu_2 \\ \Gamma/2 &= \Sigma_2^{-1} - \Sigma_1^{-1} \end{aligned}$$

Note, that with LDA we estimate Γ to be identically 0, with QDA Γ is entirely nonzero, and with ℓ_1 -PDA, Γ has both zero and nonzero elements.

3.1. Estimating Interactions. Based on the forward model above, one can consider our sparse estimation of Γ as a method for estimating sparse interactions. There has been a recent push to estimate interactions in the high dimensional setting (Radchenko and James [2010], Zhao et al. [2009], among others). The basic idea is to consider a general logistic model as in (16) (or a linear model for continuous response), and to estimate β_0 , β , and Γ in such a way that there are few nonzero entries in $\hat{\Gamma}$ (often the diagonal is constrained to be 0). The simplest of these approaches maximize a penalized logistic log-likelihood

$$\begin{aligned} \text{argmax}_{\beta, \Gamma} \quad & \sum_{i=1}^n \{y(l) \log(p_l) + (1 - y(l)) \log(1 - p_l)\} - \lambda \|\Gamma\|_1 \\ \text{s.t.} \quad & \log \left(\frac{p_l}{1 - p_l} \right) = \beta_0 + \beta^\top x_l + x_l^\top \Gamma x_l/2 \end{aligned}$$

As we have shown, for discriminant analysis considered as a forward model, nonzero off-diagonal terms in $\Gamma = \hat{\Sigma}_2^{-1} - \hat{\Sigma}_1^{-1}$ correspond to pairs of variables with interactions. Thus ℓ_1 -PDA estimates a logistic model with sparse interactions (and quadratic terms). ℓ_1 -PDA differs from other methods because it has additional distributional assumptions on the covariates which in turn put constraints on our estimates of β_0 , β , and Γ , but the underlying idea is the same.

3.2. Linear Vs Quadratic Decision Boundaries. The sparsity of Γ again shows up if we consider the decision boundaries of discriminant analysis. For each method (LDA, QDA and ℓ_1 -PDA), once the parameters are estimated, \mathbb{R}^p is partitioned into two connected spaces — one space where the estimated posterior probability of an observation is higher for class 1 and another space where it is higher for class 2. The decision boundary is $D = \{x \mid P(y = 1|x) = 0.5\}$ which is equivalent to

$\{x \mid \text{logit} [P(y = 1|x)] = 0\}$. Referring back to our forward generative framework, (14), we see that

$$D = \left\{ x \mid \hat{\beta}_0 + \hat{\beta}^\top x + x^\top \hat{\Gamma} x = 0 \right\}$$

The nonzero terms in $\hat{\Gamma} = \hat{\Sigma}_2^{-1} - \hat{\Sigma}_1^{-1}$ correspond to pairs of dimensions in which the decision boundary is quadratic rather than linear. As expected, LDA has a linear decision boundary, and QDA has a quadratic decision boundary (with all cross terms included). ℓ_1 -PDA is a hybrid of these — it is quadratic in some terms and linear in others.

4. COMPARISONS

A number of other methods have been proposed for discriminant analysis using sparsity and pooling. These methods are useful, but fill a different role than ℓ_1 -PDA. We will compare 2 of these ideas to ℓ_1 -PDA and discuss when each is appropriate.

4.1. RDA. Regularized Discriminant Analysis [Friedman, 1989] estimates the within class covariance matrices as a convex combination of the LDA and QDA estimates. Like ℓ_1 -PDA it gives a path from LDA to QDA. In contrast RDA is basis equivariant (changing the basis on which you train will not change the predictions), while ℓ_1 -PDA is not. In RDA, one uses a common idea in empirical bayes and stein estimation — we often overestimate the magnitude of extreme effects, in our case we overestimate the extremity of largest and smallest eigenvalues of $\Sigma_1 - \Sigma_2$, so RDA shrinks these values. On the other hand, ℓ_1 -PDA is very basis specific. In ℓ_1 -PDA, as in all sparse signal processing, we believe we have a good basis (in our case, we believe that the differences are sparse in this basis) and would like to leverage this in our estimation.

4.2. Sparse LDA. A number of methods have been proposed for “sparse LDA.” (Dudoit et al. [2002], Bickel and Levina [2004], Witten and Tibshirani [2011], among others). These methods either assume diagonal covariance matrices and look for sparse mean differences, or assume $\Sigma_1 = \Sigma_2$ and (either implicitly or explicitly) look for sparsity in $\Sigma^{-1}(\mu_1 - \mu_2)$. This gives a linear decision rule which uses only few of the variables. These methods are well suited to very high dimensional problems (they require many fewer observations than LDA).

In contrast ℓ_1 -PDA does not remove variables — it only shrinks decision boundaries from quadratic to linear. It is not well suited to very high dimensional problems. In particular, the solution is degenerate if

$p > n_1 + n_2$, but it will generally perform better than sparse LDA for $p < n_1 + n_2$.

To draw another parallel to logistic regression (as in Section 3.1), Sparse LDA is similar to sparse estimation of main effects (with no interactions), while ℓ_1 -PDA is similar to sparse estimation of interactions (with all main effects included).

5. OPTIMIZATION

One of the main attractions of this criterion is that it is a convex problem and hence a global optimum can be found relatively quickly. In particular we have developed a method which can solve this for up to several hundred variables (though the accuracy in poorly conditioned larger problems can be an issue).

First, for ease of notation we introduce new variables: let $A = \Sigma_1^{-1}$, $B = \Sigma_2^{-1}$, $S_A = S_1$, and $S_B = S_2$. If we plug in the sample means for μ_1 and μ_2 , our new criterion (negated for convenience) is now

$$(17) \quad \min_{A,B} -n_1 \log \det A + n_1 \operatorname{tr}(AS_A) - n_2 \log \det B$$

$$(18) \quad + n_2 \operatorname{tr}(BS_B) + \lambda \|A - B\|_1$$

subject to A, B PSD, where n_1 is the number of observations in group 1, n_2 is the number of observations in group 2. Recall that this is convex in A and B .

One could solve this using interior point methods discussed in Boyd and Vandenberghe [2004]. Unfortunately, for semi-definite programs the complexity of interior point algorithms scales like p^6 , making this approach impractical for p larger than 15 or 20. Instead we develop an approach based on the alternating direction method of moments (ADMM) which scales up to several hundred covariates.

5.1. ADMM Algorithm. ADMM is an older class of algorithms which has recently seen a re-emergence largely thanks to Boyd et al. [2010]. Our particular algorithm is a adaptation of their ADMM algorithm for sparse inverse covariance estimation. The motivation for this algorithm is simple — the combination of a logdet term and a $\|\cdot\|_1$ term makes our optimization difficult, so we split the 2 up and introduce an auxiliary variable $C \equiv A - B$ and a dual variable Γ . We leave the details of developing this algorithm to the appendix (though they are straightforward). The exact algorithm is

- (1) Initialize A_0, B_0, C_0 , and Γ_0 and choose a fixed $\rho > 0$

(2) Iterate until convergence

(a) Update A by

$$A_{k+1} = U\tilde{A}U^\top$$

where $\rho(C_k + B_k + \Gamma_k) - n_1 S_A = UDU^\top$ is its eigenvalue decomposition (with $D = \text{diag}(d_i)$), and \tilde{A} is diagonal with

$$\tilde{A}_{ii} = \frac{d_i + \sqrt{d_i^2 + 4\rho n_1}}{2\rho}$$

(b) Update B by

$$B_{k+1} = V\tilde{B}V^\top$$

where $\rho(A_{k+1} - C_k - \Gamma_k) - n_2 S_B = VEV^\top$ is its eigenvalue decomposition (with $E = \text{diag}(e_i)$) and \tilde{B} is diagonal with

$$\tilde{B}_{ii} = \frac{e_i + \sqrt{e_i^2 + 4\rho n_2}}{2\rho}$$

(c) Update C by

$$C_{k+1} = S_{\lambda/\rho}(A_{k+1} - B_{k+1} - \Gamma_k)$$

where $S_\lambda(\cdot)$ is the element-wise soft thresholding operator

$$S_\lambda(Z)_{i,j} = \text{sign}(Z_{i,j}) \max(|Z_{i,j}| - \lambda, 0)$$

(d) update Γ by

$$\Gamma_{k+1} = \Gamma_k + \rho(C_{k+1} - A_{k+1} + B_{k+1})$$

Upon convergence, A^* and B^* are the variables of interest (the rest may be discarded). The complexity of each step of this algorithm is dominated by the eigenvalue decompositions, each of which require $O(p^3)$ operations.

6. PATH-WISE SOLUTION

Often we do not know a-priori what value our regularization parameter should be and would like to fit the entire path from λ_{max} (corresponding to the LDA solution) to $\lambda = 0$ (corresponding to the QDA solution). We define

$$\lambda_{max} \equiv \frac{n_1 n_2 \|S_1 - S_2\|_\infty}{n_1 + n_2}$$

It is easy to see that for $\lambda \geq \lambda_{max}$, $\hat{\Sigma}_1 = \hat{\Sigma}_2 = \frac{n_1 S_1 + n_2 S_2}{n_1 + n_2}$ (our LDA solution) satisfies (10) and (11), and thus is our solution. One can also see that $\hat{\Gamma} = \frac{n_1 n_2 (S_1 + S_2)}{n_1 + n_2}$ is our optimal dual variable for $\lambda \geq \lambda_{max}$.

To solve along a path we start at $\lambda = \lambda_{max}$, and plug in our known solution. We then decrease λ and solve the new problem, initializing our algorithm at the previous $\hat{\Sigma}_1$, $\hat{\Sigma}_2$, and $\hat{\Gamma}$. Because λ changes only slightly (and thus our solution changes only slightly), this approach is very efficient as compared to solving from scratch at each λ . When S_A and S_B are full rank our QDA solution is well defined and it is possible to run our path all the way to $\lambda = 0$. Due to convergence issues along the potentially poorly conditioned end of the path (which we discuss in the next section) we instead choose to set $\lambda_{min} = \epsilon \lambda_{max}$ and log-linearly interpolate between the two (in our implementation default ϵ value is 0.01).

6.1. Convergence Issues. While ADMM is a good algorithm for finding an near exact solution, it is not considered a great algorithm for an exact solution (though it does eventually converge to arbitrary tolerance, this may require an unwieldy number of iterations). In our application, solving to machine tolerance is unnecessary (the statistical uncertainties are much greater than this). However, in some cases (especially with $p \sim n_1 + n_2$), near the end of the path our solution converges extremely slowly. Unfortunately there is no simple fix for this (more accurate interior point algorithms don't scale beyond 15 or 20 variables). While not ideal, this does not overly concern us — convergence is slow in the region where $\Sigma_1^{-1} - \Sigma_2^{-1}$ is not very sparse (a region where we expect ℓ_1 -PDA to perform poorly anyways). We will see an example of this issue arise later in Section 8.

One should also note that convergence rates near the end of the path are highly dependent on our choice of ρ . This is characteristic of all ADMM problems. To date, finding a disciplined choice of ρ for ADMM is still an open question. We use $\rho = 1$ as our default, as it appears to work reasonably well for a range of problems.

7. SIMULATED AND REAL DATA

To show its efficacy, we applied ℓ_1 -PDA to real and simulated data. In both cases we compare our method to linear, quadratic and regularized discriminant analysis and show improvement over both in overall classification error and on ROC plots. In all problems ℓ_1 -PDA was fit with 30 lambda values log-linearly interpolating λ_{max} and $0.01 * \lambda_{max}$. RDA was fit with 30 equally spaced λ -values between 0 and 1.

7.1. Simulated Data. We simulated data from the two class gaussian model described in Section 1 with $p = 30$ covariates. We set $\Sigma_1 = I_{p \times p}$




FIGURE 1. Average ROC curve for simulated data with $n_k = 33$, $c = 0.9$

and

$$\Sigma_2 = \left(\begin{array}{c|c} \text{---} \frac{C}{0} \text{---} & \text{---} 0 \text{---} \\ \hline \text{---} 0 \text{---} & I_{(p-5) \times (p-5)} \end{array} \right)$$

where C is 5×5 matrix with constant value c on the off diagonal entries, and 1 on the diagonal. We also set a small mean difference between the groups: $\mu_1 = \underline{0}_p$

$$\mu_2 = \left(\begin{array}{c} \text{---} \frac{\Delta}{\underline{0}_{(p-10)}} \text{---} \end{array} \right)$$

where Δ is a 10-vector of 1s

Under this model $\Sigma_1^{-1} - \Sigma_2^{-1}$ is nonzero only in the upper left 5×5 submatrix. We simulated using varying numbers of observations $n_1 = n_2 \in (33, 40, 60)$, and values of $c \in (0.3, 0.6, 0.9)$. We used 3 data sets for each simulation — one to fit the initial model, one to choose the optimal value of λ and our final set to get an unbiased estimate of misclassification error.

As you can see from Table 7.1, when the signal to noise ratio (SNR) is too small ℓ_1 -PDA adaptively shrinks towards LDA and sees similar performance. When SNR is sufficiently large (the third group in the table), ℓ_1 -PDA is able to pick out the nonzero entries and achieves substantially better misclassification rates. In these cases RDA also does fairly well (adaptively choosing between LDA and QDA), however because it

		# Observations per Group (n_k)		
		33	40	60
$c = 0.3$	ℓ_1 -PDA	0.82	0.85	0.88
	LDA	0.82	0.85	0.88
	QDA	0.58	0.65	0.74
	RDA	0.81	0.85	0.88
$c = 0.6$	ℓ_1 -PDA	0.82	0.83	0.87
	LDA	0.82	0.84	0.86
	QDA	0.59	0.66	0.76
	RDA	0.81	0.83	0.86
$c = 0.9$	ℓ_1 -PDA	0.84	0.88	0.92
	LDA	0.80	0.83	0.85
	QDA	0.65	0.76	0.86
	RDA	0.81	0.84	0.88

TABLE 1. Average % of correct classifications over 100 simulated datasets (standard errors for all entries are less than 0.006)

does not take sparsity into account, it is outperformed by ℓ_1 -PDA. We consider the large SNR case more carefully in Figure 1 (an ROC curve for $n_k = 33$, $c = 0.9$) Again we used 3 data sets per realization to get an unbiased curve estimate (and ran 100 random realizations, though only average is shown on Figure 1). We estimated AUC for each procedure: ℓ_1 -PDA 0.924 ± 0.002 , LDA 0.875 ± 0.003 , QDA 0.732 ± 0.007 , and RDA 0.887 ± 0.003 . ℓ_1 -PDA does substantially better than LDA, QDA, and RDA. With $p = 30$ and $n_k = 33$ there is clearly not enough data for QDA to perform well (though the sample correlation matrices are still invertible). However, as noted, ℓ_1 -PDA also has a large edge over LDA and RDA.




FIGURE 2. Plot of validated prediction accuracy for regularization path in 58 mines/rocks, with $\lambda_{\min} = 0.01\lambda_{\max}$ for ℓ_1 -PDA, and ROC curve for λ_{11}

8. REAL DATA

We also applied ℓ_1 -PDA to the “Sonar, Mines vs. Rocks” data [Gorman and Sejnowski, 2010]. This dataset has 60 sonar signals measured on each of 208 objects (each labeled as either a rock or a mine). We randomly chose 150 Mines/Rocks to train with, and then classified the remaining 58.

As one can see from Figure 2, ℓ_1 -PDA performs better on this data than either LDA, QDA or RDA. Estimated true classification rate peaks near the middle of our regularization path, showing that a fair amount of regularization can significantly improve classification. As we mentioned in Section 6.1 one can see convergence issues near the end of our path — we would expect the CV error at our 30th λ -value to nearly match that of QDA (nearly rather than exactly because we don’t run to $\lambda_{\min} = 0$). However, it does not, indicating that our solution is not converging to the QDA solution. This does not overly concern us as our validation error reaches its crest well before this.

We also see an ROC curve comparing ℓ_1 -PDA, LDA, QDA, and RDA. For RDA we chose the simplest model which maximized predictive accuracy (the 21st λ value), and for ℓ_1 -PDA the tenth λ value, the most regularized model before a precipitous drop in predictive accuracy (so as to minimize bias for ℓ_1 -PDA). The ℓ_1 -PDA curve may still be slightly

biased as we chose it from a section of our path seen to do well in overall classification error (though not the peak). Nonetheless, this curve appears indicative of an advantage from ℓ_1 -PDA over LDA, QDA, and RDA.

9. DISCUSSION

In this paper we proposed ℓ_1 -PDA, a classification method for gaussian data which adaptively pools the precision matrices. We motivated our method, and showed connections between it and estimating sparse interactions. We gave two efficient algorithms to fit have ℓ_1 -PDA, and have shown its efficacy on real and simulated data. We have made and plan to provide an R implementation for ℓ_1 -PDA publically available on CRAN.

10. APPENDIX α

We include a short overview of the ADMM algorithm. We can rewrite (17) as

$$\begin{aligned} \min_{A,B} & -n_1 \log \det A + n_1 \operatorname{tr}(AS_A) - n_2 \log \det B \\ & + n_2 \operatorname{tr}(BS_B) + \lambda \|C\|_1 \\ \text{s.t. } & C = A - B \end{aligned}$$

At the optimum we have $C = A - B$, so this is equivalent to

$$(19) \quad \min_{A,B,C} -n_1 \log \det A + n_1 \operatorname{tr}(AS_A) - n_2 \log \det B$$

$$(20) \quad + n_2 \operatorname{tr}(BS_B) + \lambda \|C\|_1 + \frac{\rho}{2} \|C - A + B\|_F^2$$

$$(21) \quad \text{s.t. } C = A - B$$

ρ can be any fixed positive number (though its choice will affect the convergence rate of algorithm). We will motivate this addition shortly. Now, using strong duality, we can move our constraint into the objective, and finally arrive at

$$(22)$$

$$\begin{aligned} \max_{\Gamma} \min_{A,B,C} & -n_1 \log \det A + n_1 \operatorname{tr}(AS_A) - n_2 \log \det B \\ (23) \quad & + n_2 \operatorname{tr}(BS_B) + \lambda \|C\|_1 + \rho \operatorname{trace} [\Gamma^\top (C - A + B)] \end{aligned}$$

$$(24) \quad + \frac{\rho}{2} \|C - A + B\|_F^2$$

For ease of notation we denote

$$\begin{aligned}\psi_\Gamma(A, B, C) = & -n_1 \log \det A + n_1 \operatorname{tr}(AS_A) - n_2 \log \det B \\ & + n_2 \operatorname{tr}(BS_B) + \lambda \|C\|_1 + \rho \operatorname{trace} [\Gamma^\top (C - A + B)] \\ & + \frac{\rho}{2} \|C - (A - B)\|_F^2\end{aligned}$$

and

$$M(\Gamma) = \min_{A, B, C} \psi_\Gamma(A, B, C).$$

Now, by basic convex analysis, the dual of any strongly convex function (with convexity constant ρ) is differentiable and its derivative has lipschitz constant ρ . Unfortunately (19) is not necessarily strongly convex, however the addition of $\|C - A + B\|_F^2$, affords it many of the same properties. In particular if C^* , A^* , B^* are the argmin of ψ_{Γ_0} for a given Γ_0 , then

$$\left. \frac{\partial}{\partial \Gamma} M \right|_{\Gamma_0} = C^* - A^* + B^*$$

If we could easily calculate $M(\Gamma)$, then we could use gradient ascent on Γ

$$\Gamma_{k+1} = \Gamma_k + \rho(C_k^* - A_k^* + B_k^*)$$

and one would have A_k^* and B_k^* converging to the argmax of our original problem (4). Unfortunately, $M(\Gamma)$ is not easy to calculate, however ψ_Γ is relatively simple to minimize in one variable at a time (A , B , or C) with all other variables fixed. In ADMM we employ the same idea as gradient descent, only we fudge the details — instead of actually calculating $M(\Gamma)$, we minimize first in A , with B , and C fixed, then in B with A and C fixed and finally in C with A and B fixed. After these 3 updates, we take our “gradient” step as before (though this time it is not a true gradient step). This leads to the following algorithm:

- (1) Initialize A_0 , B_0 , C_0 , and Γ_0
- (2) Iterate until convergence
 - (a) Update A by

$$A_{k+1} = \operatorname{argmin}_A \psi_{\Gamma_k}(A, B_k, C_k)$$

- (b) Update B by

$$B_{k+1} = \operatorname{argmin}_B \psi_{\Gamma_k}(A_{k+1}, B, C_k)$$

- (c) Update C by

$$C_{k+1} = \operatorname{argmin}_C \psi_{\Gamma_k}(A_{k+1}, B_{k+1}, C)$$

- (d) Take “gradient step”; update Γ by

$$\Gamma_{k+1} = \Gamma_k + \rho(C_{k+1} - A_{k+1} + B_{k+1})$$

One may note that if we instead iterate steps $a - c$ to convergence each time before taking step d , we end up again with gradient descent.

10.1. Inner Loop Updates. In this section we derive the exact updates for A , B , and C in steps a , b and c of our ADMM algorithm. We begin with A : to find $\operatorname{argmin}_A \psi_{\Gamma_k}(A, B_k, C_k)$ we must minimize

$$\begin{aligned} & -n_1 \log \det A + n_1 \operatorname{tr}(AS_A) + \rho \operatorname{trace} [\Gamma_k^\top (C_k - A + B_k)] \\ & + \frac{\rho}{2} \|C_k - A + B_k\|_F^2 \end{aligned}$$

If we take the derivative of this and set it equal to 0 we get

$$(25) \quad \rho A - n_1 A^{-1} = \rho (C_k + B_k + \Gamma_k) - n_1 S_A$$

Now if we let $\rho (C_k + B_k + \Gamma_k) - n_1 S_A = UDU^\top$ be its eigenvalue decomposition (with $D = \operatorname{diag}(d_i)$), then (25) is satisfied by

$$A = U \tilde{A} U^\top$$

where \tilde{A} is diagonal and

$$\tilde{A}_{ii} = \frac{d_i + \sqrt{d_i^2 + 4\rho n_1}}{2\rho}$$

We can solve for B_{k+1} similarly. Let $\rho (A_{k+1} - C_k - \Gamma_k) - n_2 S_B = VEV^\top$ be its eigenvalue decomposition (with $E = \operatorname{diag}(e_i)$). Then $\operatorname{argmin}_B \psi_{\Gamma_k}(A_{k+1}, B_k, C_k)$ is

$$B = V \tilde{B} V^\top$$

where \tilde{B} is diagonal and

$$\tilde{B}_{ii} = \frac{e_i + \sqrt{e_i^2 + 4\rho n_2}}{2\rho}$$

The last variable to solve for is C . Ignoring all terms without a C , we need to minimize

$$\lambda \|C\|_1 + \rho \operatorname{trace} [\Gamma^\top (C - A + B)] + \frac{\rho}{2} \|C - A + B\|_F^2$$

This is equivalent to minimizing

$$\frac{1}{2} \|C - (A_{k+1} - B_{k+1} - \Gamma)\|_F^2 + \frac{\lambda}{\rho} \|C\|_1$$

which is solved by

$$C = S_{\lambda/\rho} (A_{k+1} - B_{k+1} - \Gamma)$$

where $S_{\lambda/\rho}$ is the entry-wise soft thresholding operator on the entries of the matrix. For $i \neq j$

$$S_{\lambda/\rho}(X)_{ij} = \operatorname{sign}(X_{ij}) \max(|X_{ij}| - \lambda/\rho, 0)$$

So, in full detail, our algorithm is

- (1) Initialize A_0 , B_0 , C_0 , and Γ_0
- (2) Iterate until convergence
 - (a) Update A by

$$A_{k+1} = U\tilde{A}U^\top$$

where $\rho(C_k - B_k + \Gamma_k) - n_1 S_A = UDU^\top$ is its eigenvalue decomposition (with $D = \text{diag}(d_i)$), and \tilde{A} is diagonal with

$$\tilde{A}_{ii} = \frac{d_i + \sqrt{d_i^2 + 4\rho n_1}}{2\rho}$$

- (b) Update B by

$$B_{k+1} = V\tilde{B}V^\top$$

where $\rho(A_{k+1} - C_k - \Gamma_k) - n_2 S_B = VEV^\top$ is its eigenvalue decomposition (with $E = \text{diag}(e_i)$) and \tilde{B} is diagonal with

$$\tilde{B}_{ii} = \frac{e_i + \sqrt{e_i^2 + 4\rho n_2}}{2\rho}$$

- (c) Update C by

$$C_{k+1} = S_{\lambda/\rho}(A_{k+1} - B_{k+1} - \Gamma_k)$$

- (d) update Γ by

$$\Gamma_{k+1} = \Gamma_k + \rho(C_{k+1} - A_{k+1} + B_{k+1})$$

The complexity of each step of this algorithm is dominated by the eigenvalue decompositions, each of which require $O(p^3)$ operations. For this reason, while the algorithm can solve problems for p in the hundreds, it will be difficult to scale to larger problems. One should note that p in the hundreds is already an optimization problem with tens of thousands of variables.

REFERENCES

- O. Banerjee, L. E. Ghaoui, and A. d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research*, 9: 485–516, 2008.
- P. Bickel and E. Levina. Some theory for fisher's linear discriminant function, 'naive bayes', and some alternatives when there are many more variables than observations. *Bernoulli*, pages 989–1010, 2004.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Machine Learning*, 3(1):1–123, 2010.
- S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, pages 33–61, 1998.
- P. Danaher, P. Wang, and D. Witten. The joint graphical lasso for inverse covariance estimation across multiple classes. *Arxiv preprint arXiv:1111.0324*, 2011.
- S. Dudoit, J. Fridlyand, and T. Speed. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American statistical association*, 97(457):77–87, 2002.
- J. Friedman. Regularized discriminant analysis. *Journal of the American statistical association*, pages 165–175, 1989.
- R. Gorman and T. Sejnowski. Uci: Machine learning repository, 2010. URL <http://archive.ics.uci.edu/ml>.
- P. Radchenko and G. James. Variable selection using adaptive nonlinear interaction structures in high dimensions. *Journal of the American Statistical Association*, 105(492):1541–1553, 2010.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society B*, 58:267–288, 1996.
- D. Witten and R. Tibshirani. Penalized classification using fisher’s linear discriminant. *Journal of the Royal Statistical Society, Series B*, 2011.
- P. Zhao, G. Rocha, and B. Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *The Annals of Statistics*, 37(6A):3468–3497, 2009.